


Jonathan Gleason

DATA MANAGEMENT STATUS

Data Management Team
CERES Science Team Meeting October 4, 2011




Overview



- Organization
- Activity Since Last STM
- Data Availability
- Platform Migration Status
- NPP Testing & Preparation
- AMI-P Stress Test Results



Data Management Team Members



Instrument:

Denise Cooper
Thomas Grepiotis
Richard Spivak
Mark Timcoe
Nelson Hillyer
Dianne Snyder

ERBElike:

Dale Walikainen
Jeremy Lande

TISA Gridding:

Raja Raju
Hazari Syed

Clouds:

Sunny Sun-Mack
Ricky Brown
Yan Chen
Liz Heckert
Rita Smith
Sharon Gibson

Convolution:

Walter Miller
Igor Antropov

TISA Averaging:

Cathy Nguyen
Betty Lock

Production / Optimization:

Lisa Coleman
Carla Grune

Inversion:

Victor Sothcott

SARB:

Tom Caldwell

Configuration Management:

Tammy Ayers
Joanne Saunders

SSAI:

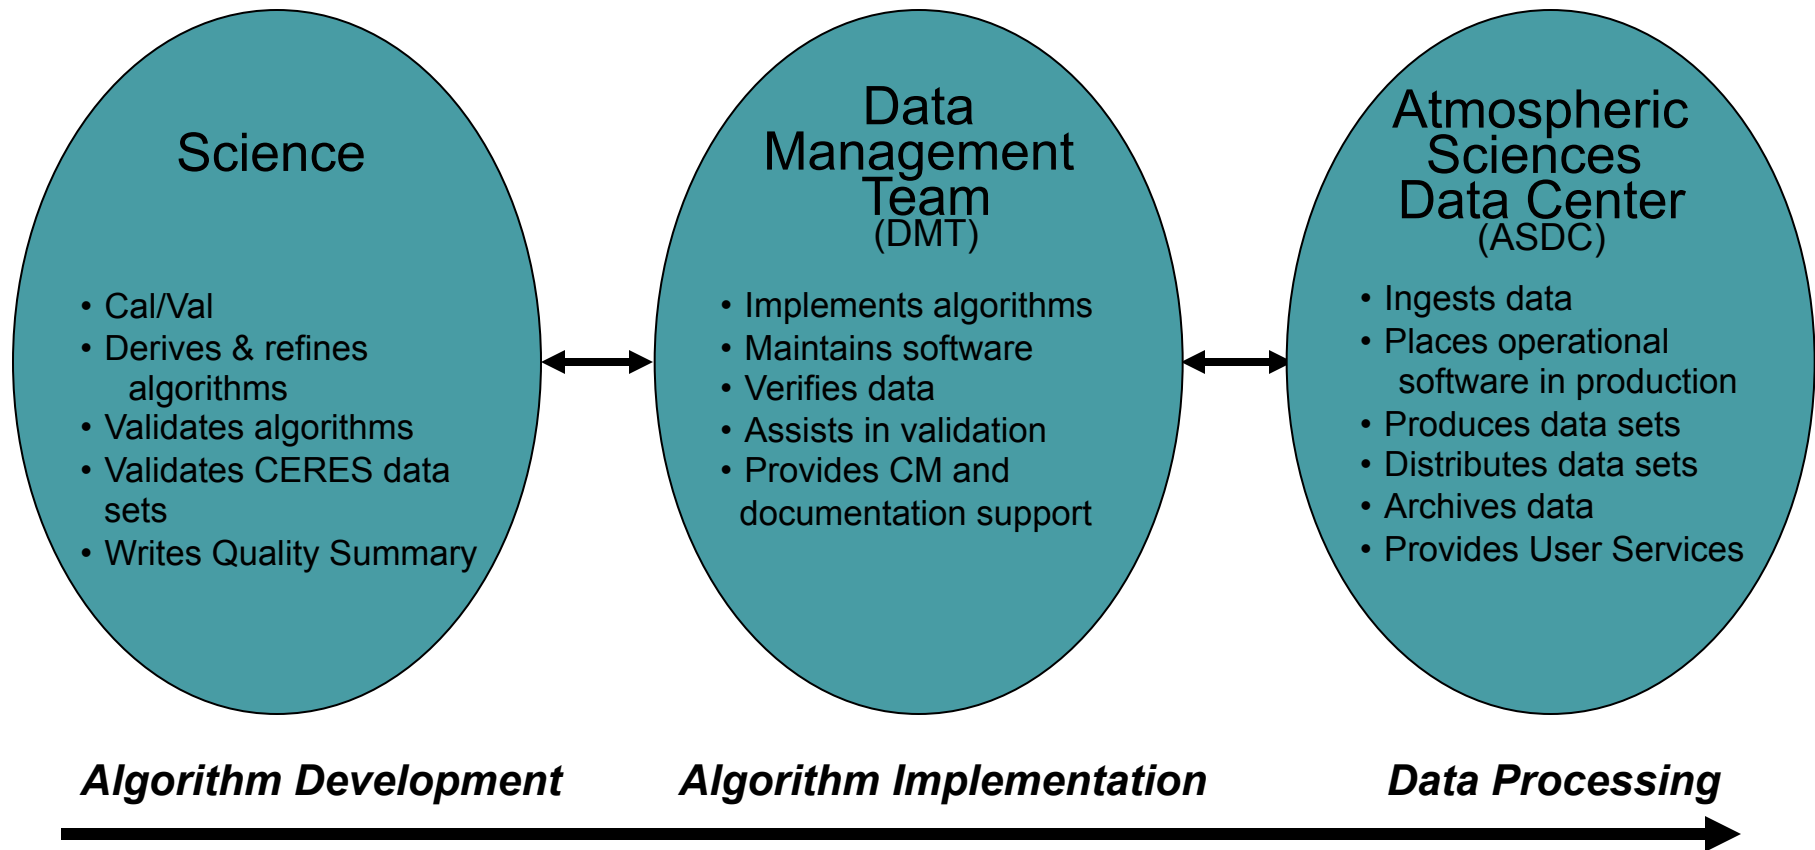

Lee Bodden

FM5:

Jim Closs




CERES Organization





CERES Subsystems




- CERES is made up of 7 Working Groups
 - Instrument
 - ERBElike
 - Clouds
 - Inversion or ADM
 - SOFA
 - SARB
 - TISA
- Code organized into 12 Subsystems
 - Each subsystem tied to 1 or more working groups
- Each Subsystem made up of 1 or more Product Generation Executives (PGEs)
 - Currently there are 73 active PGEs

CERES Software Overview

Subsystem Number	Subsystem Name	Number of PGEs	LOC (to nearest 1K)	Publicly Available Data Products	Product Frequency
1	Instrument/Pre-Processor	1	4K		
1	Instrument	11	285K	BDS	1/day
2	ERBE-like/ Inversion	5	23K	ES-8	1/day
3	ERBE-like/ TSA	2	12K	ES-9, ES-4	1/month
4.1 – 4.4	Clouds/VIIRS Subset Code	1	21K	VIMD	12/hour
4.1 – 4.4	Clouds	10	359K		
4.5 – 4.6	Inversion	13	164K	SSF	1/hour
5	SARB	6	164K	CRS	1/hour
6 & 9	TISA-Gridding	13	72K	CRS1deg-Hour, SSF1deg-Hour, ISCCP-D2like-Day/Nit	60/month, 36/month, 1/month
11	GGEO	5	170K	ISCCP-D2like-GEO	1/month
7.2	Synoptic SARB	1	47K		
7.1 & 8 10	TISA-Averaging	3	87K	SSF1deg-Day, SSF1deg-Month, SYN1deg-(3Hour, M3hour, Mhour, Month)	1/day, 1/ month, 1/ month 5/month
12	MOA	2	14K		
	CERESlib		145K		
		Total LOC	1,548K		



Activity since Spring STM



25 total deliveries since 4/26/2011

- Instrument (12)
 - 2 routine (1.4P1 & 1.4P2), 4 bug fixes found during NPP testing (TVAC, NCT3 and NCT4), 4 bug fixes found during internal validation, 1 FM5 ancillary files ,1 Ada risk mitigation
- ERBE-Like (1)
 - All PGEs to AMI-P
- Clouds (1)
 - 1 Beta Edition 4
- Tisa Grid (2)
 - 1 ISCCP-D2Like Day/Nit, 1 Ed3 xtk LUTs
- Tisa Averaging (2)
 - 1 New Ed3 TSI, 1 TSI ancillary
- Inst SARB (1)
 - 1 MATCH files
- GGEO (2)
 - 2 Ancillary input files
- MOA (2)
 - 1 Ed 4 to AMI-P, 1 Script update
- CERESLib (1)
 - 1 metadata support for NPP and merged Terra-Aqua products
- Perl_Lib (1)
 - 1 added routines




Data Availability



Product	Instruments	Available through	Comments
BDS	Terra & Aqua	Jan 1, 2011	Need Input
SSF	Terra	3/1/2000 – 7/31/2003 & 10/31/2005 – 12/31/2010	Currently Processing
	Aqua	4/1/2005 – 12/31/2010	
SFC	Terra	Nov 2005 – Nov 2010	
	Aqua	April 2005 – Nov 2010	
SSF Edition 1-CV	Terra Aqua	7/1/2010 – 7/31/2011	Internal Product
ISCCP-D2like-Day ISCCP-D2like-Nite		Jan 2002 – Dec 2007	Currently Processing
ISCCP-D2like-GEO		Mar 2000 – Feb 2010	



Computing Platform Migration





- Instrument production on AMI-P Mid October

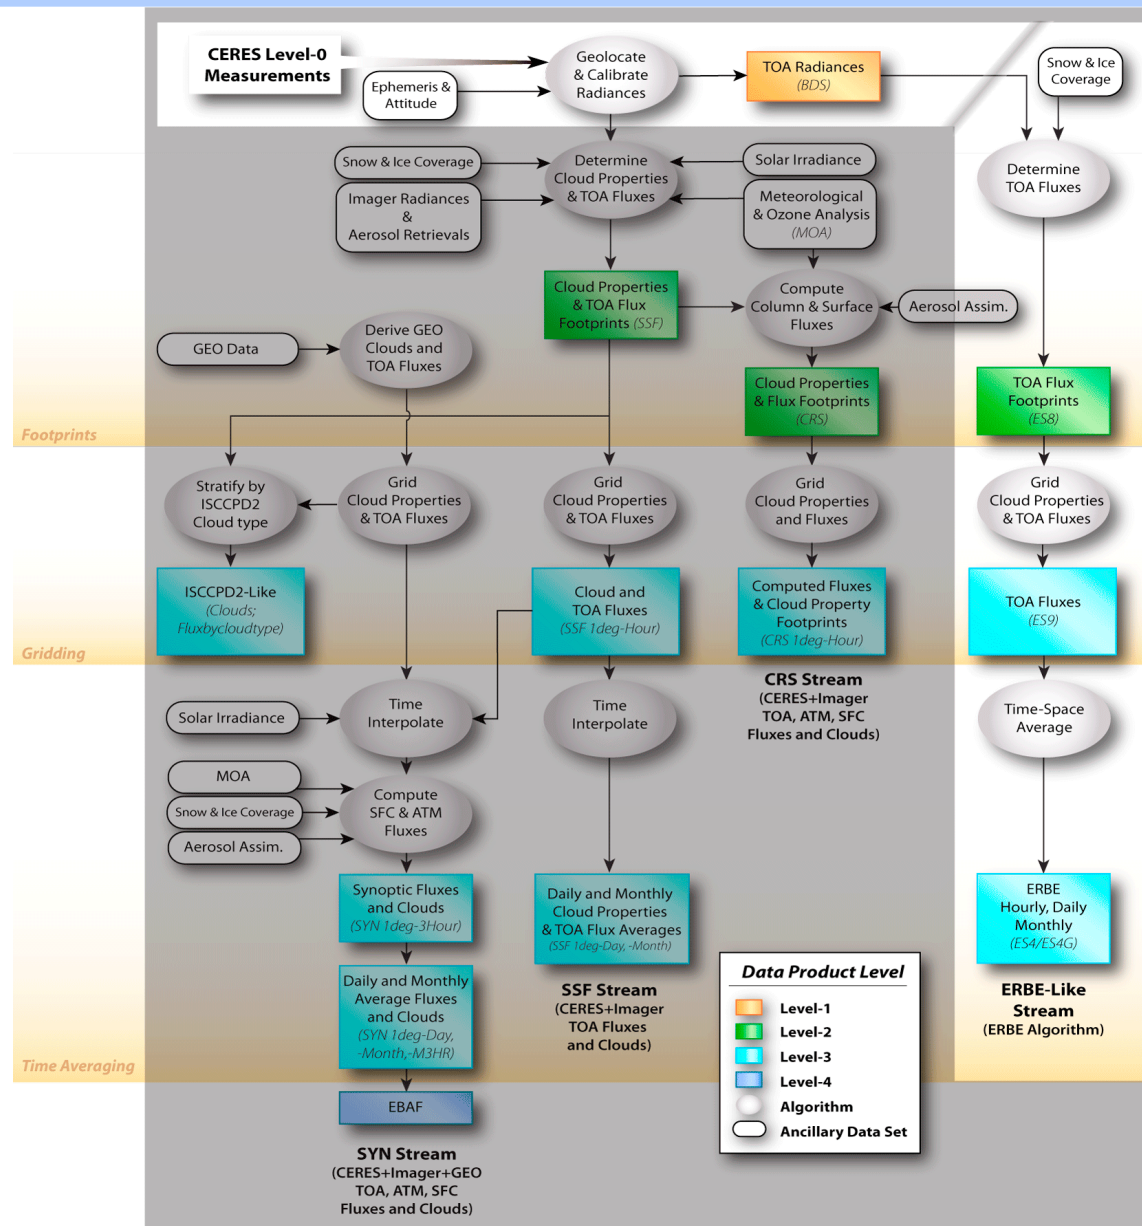
Warlock (SGI)	Magneto (P4)	AMI-P (P6 & x86)
<ul style="list-style-type: none">• Instrument• FLASHflux	<ul style="list-style-type: none">• ERBElike• Clouds Ed1-CV• Inversion Ed2 & Ed3• TISA Gridding Ed3• Inst. SARB Ed2• Synoptic SARB Ed3• GGEO• MOA Ed2	<ul style="list-style-type: none">• Instrument• MOA Ed4• ISCCP-D2like Day/Nit + GEO• FLASHflux <p>In Testing</p> <ul style="list-style-type: none">• ERBE-Like• Clouds Ed4*• Synoptic SARB Ed3• TISA Averaging Ed3* <p>In Development</p> <ul style="list-style-type: none">• Inversion Ed4



Production Schedule




NPP Code to Run at Launch





New Software for FM5



- Instrument Preprocessor: Converts NPP RDR format to EOS Level 0 format - Extracts spacecraft diary and radiance data from NPP RDR format and creates ephemeris, attitude, radiance files
- Instrument Subsystem Ada to C++: Converted 9 Ada PGEs to 3 C++ PGEs (1 remaining to convert)
- Created 1 new Ada PGE to handle FM5 functionality as risk mitigation for C++
- VIIRS radiance and geolocation subset code (runs at Land PEATE)



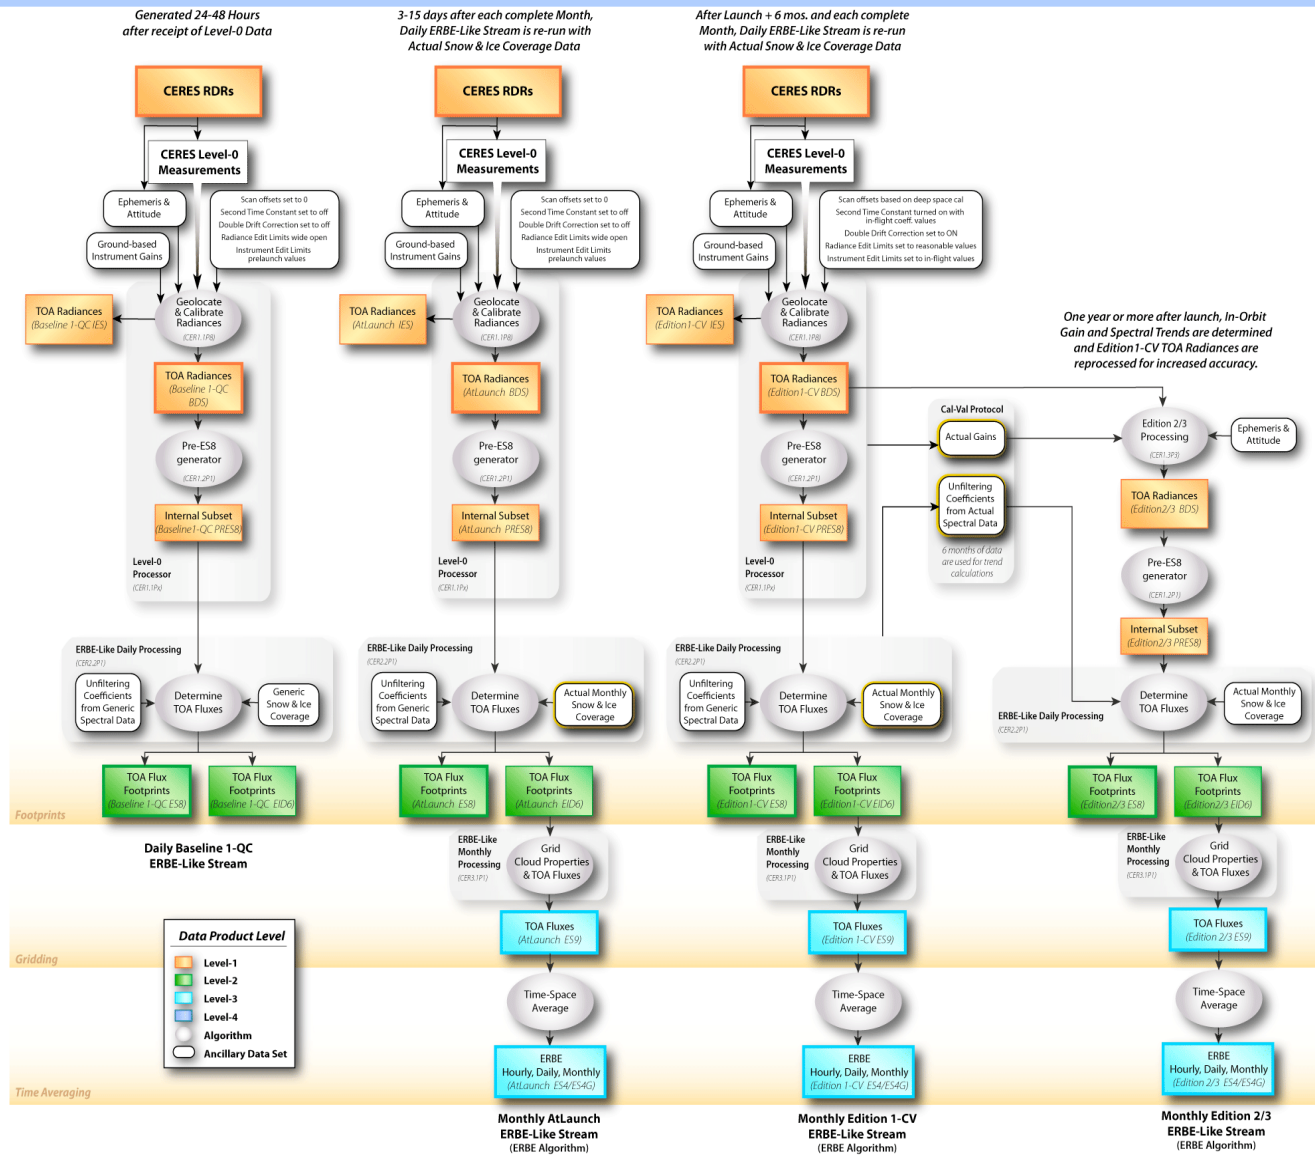

CERES NPP PGE Configuration



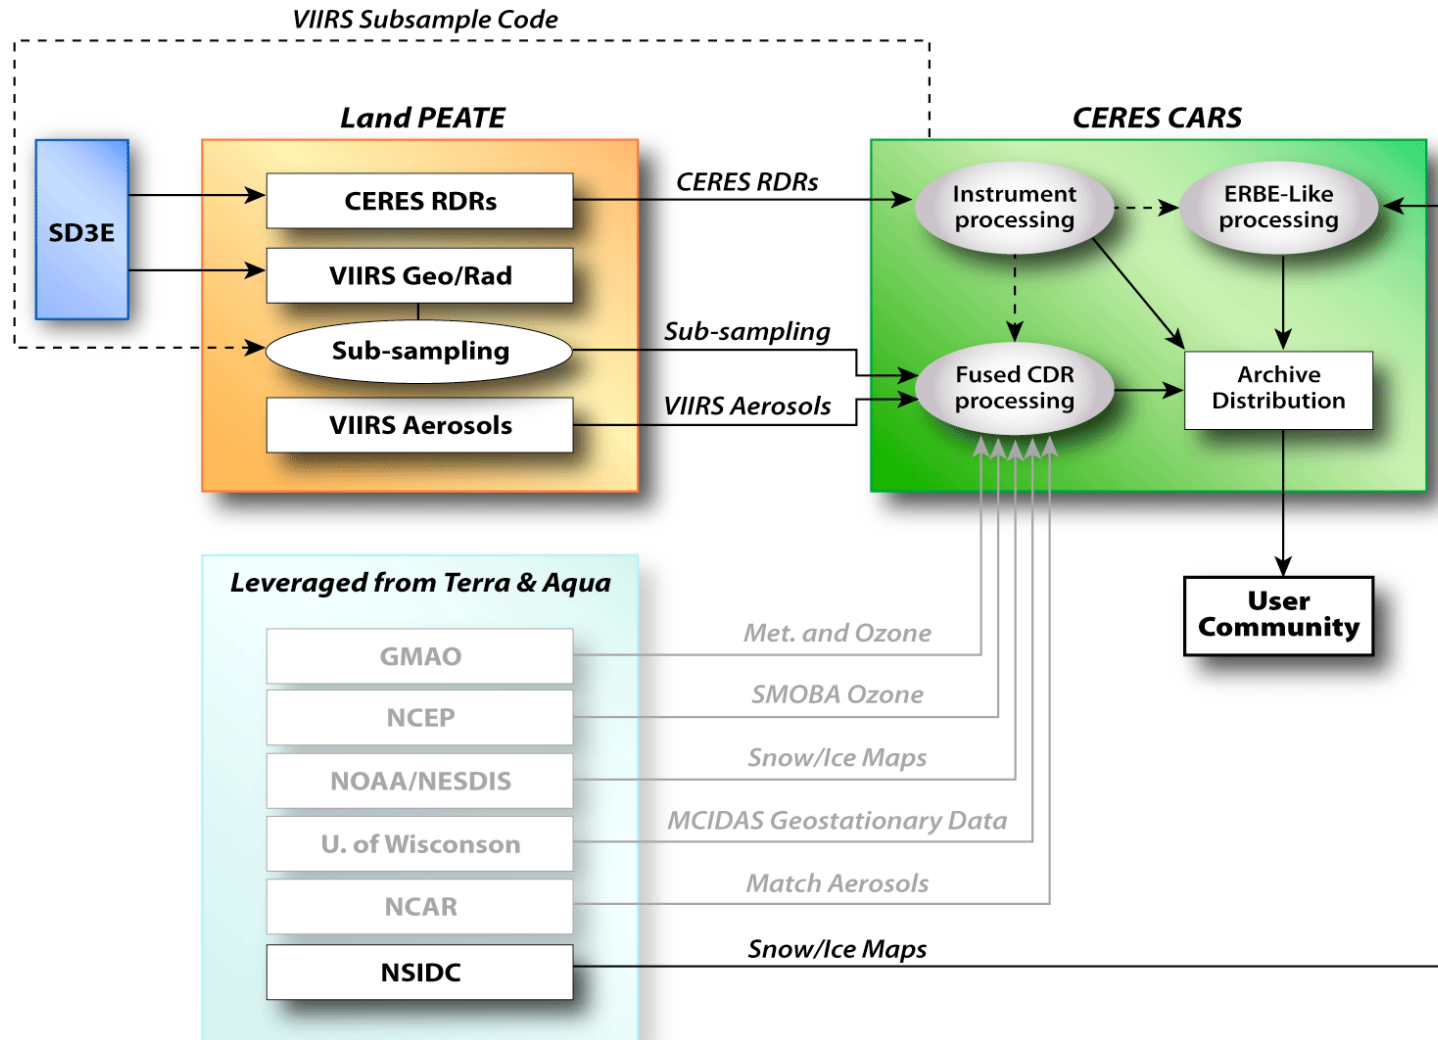
PGE	Description	Language	Current Platform	Launch Platform	Launch Tested	Final Version
1.0P2	Preprocessor	C++	x86 +P6	x86 + P6	Yes	x86 + P6
1.1P8	Main Processor	C++	P6	P6	Yes	x86 + P6
1.2P1	PRE Es8 Generator	C	x86 +P6	x86 + P6	Yes	x86 + P6
1.4P1	BDSI Subset	C++	x86 +P6	x86 +P6	Yes	x86 + P6
1.4P2	Gain Analyzer	C++	x86+P6	x86 +P6	Yes	x86 + P6
1.4P3	Ed2/3 Main Processor	C++	none	x86 +P6	No	x86 + P6
2.1P1	Snow/Ice map	F90	P4	P6	Yes	x86 + P6
2.4P1	SCC generator	F90	P4	P6	Yes	x86 + P6
2.2P1	ES8 Generator	F90 + C	P4	P6	Yes	x86 + P6
2.3P1	Previous 6 hrs	F90	P4	P6	Yes	x86 + P6
2.3P2	Following 6 hrs	F90	P4	P6	Yes	x86 + P6
3.1P1	ES4+ES9 Generator	F90 + C	P4	P6	Yes	x86 + P6



Processing Overview




Interfaces






NPP Testing & Preparation



- Ground System Interface Test (GSIT) – March 2010
- NPP Confidence Test 3 (NCT3) – April 2011
- NPP Confidence Test 4 (NCT4) – July/August 2011
- CERES FM5 Science Operations Readiness Review – September 2011
- Ground Operations Exercise (GOE) / Mission Dress Rehearsal – September/October 2011




NPP Testing



- Ground System Interface Test (GSIT) – March 2010
 - Proxy data sent thru SDS to Land PEATE to ASDC
 - Software Issues (1) – Delivered 1.1P8 with full functionality
- NPP Confidence Test 3 (NCT 3) – April 2011
 - Proxy data sent thru IDPS to SDS to Land PEATE to ASDC
 - Software Issues (3)
 - Preprocessor read incorrect files and used custom time conversion routines
 - Pre-ES8 generator could not handle FM5 in metadata – Corrected via CERESLib routine and recompile
 - ERBE-Like could not execute due to Pre-ES8 input – Successfully retested after CERESLib routine provided




NPP Testing



■ NPP Confidence Test 4 (NCT 4) – August 2011

- 7 days proxy (3 days repeated) & 4 days TVAC data sent thru IDPS to SDS to Land PEATE to ASDC
- Software Issues (2)
 - Instrument QC reports missing FM5 Instrument name in QC reports
 - Pre-ES8 generator incorrectly byte swapped end date/time field – ES8 could not archive as no end date/time provided in metadata
- Production scripts pointed to incorrect data input directory – wrong data ingested
- Discovered proxy data orbit spatially shifted by 1 hour 40 minutes (Antarctica showed largest TAO flux values)


LW TOA flux ES8-gif for NCT4 (9/6/10)



Note Antarctica has the largest (red) TOA fluxes




FM5 Science ORR



- 1 Day, Science Directorate, internal review September 16, 2011 (6 member review panel)
- **RFA 1:** Provide complete schedule for completion of remaining launch support PGEs
- **RFA 2:** Identify ASDC hardware system configuration freeze timeline for pre and post NPP launch period
- **RFA 3:** Provide migration plan for Instrument and ERBE-Like websites (required at launch by Instrument team) from IpoSUN and SPG10 legacy SGI systems to SD web cluster environment
- **RFA 4:** Identify migration plan with scope of work for 1.1P8 to x86 platform
- **Recommendation 1:** Implement stress test scripting approach for all CERES production on AMI-P
- **Recommendation 2:** Synchronize SCF and ASDC production environments (Compilers, Toolkits, OS versions) as per lesson learned from Data Management Team in review
- **Recommendation 3:** Implement consistent software config. Management practices across DMT and Instrument team and make transparent to other Directorate teams for lesson sharing




Conversion of Remaining PGEs to x86



- CER1.1P8 will be updated to work on either platform by Launch + 6 months.
 - All areas where byte-swapping identified in Ada Version
 - Memory leak discovered in C++ and will be addressed
- Three ERBE-Like PGEs successfully compiled on x86 platform, Four remaining to convert
 - Plot software for ES8 does not function on x86 (expected due to possible byte swapping issue)



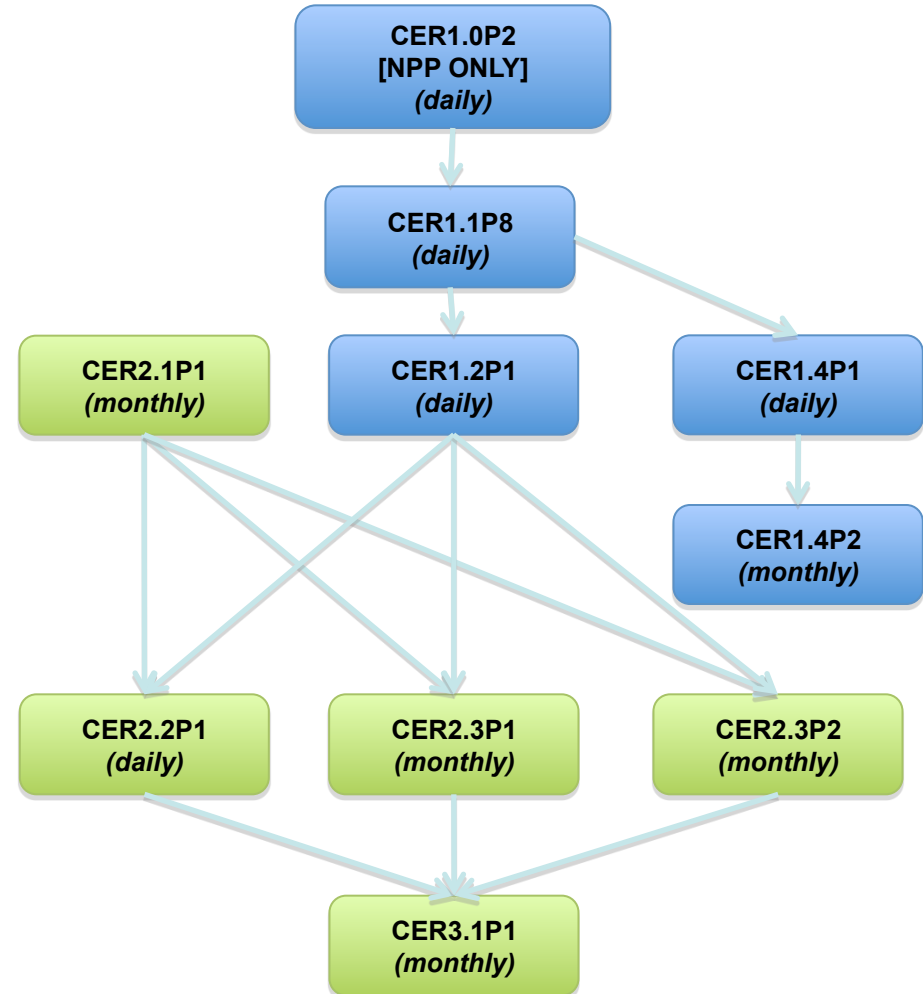
1.1P8 Terra/Aqua Validation



- 1.1P8 validation efforts for Terra and Aqua data identified critical error in code
- Error cases (corrupt input file) always exit success – Ada implicitly handles exceptions where C++ requires developer to explicitly throw and catch exceptions.
- Discovered error in how Radiance and mode flags were byte swapped (didn't see until compared with IES, not in typical validation effort)


AMI-P Stress Test

- Can AMI-P reasonably support processing for FM5 alongside FM1-4 operations?
- Exercise AMI-P processing hardware using Instrument and ERBE-like PGEs.
- Goal is to exercise all possible processing streams for Instrument and ERBE-like. End-to-end processing from CER1.0P2 to CER1.4P2 and CER3.1P1.
- **CER1.0P2** [FM5], **CER1.1P8** [FM1-4], and **CER2.1P1** [FM1-5] are immediately runnable at start of test.
- Maximize system load by queuing high volume of CERES reprocessing jobs.





AMI-P Stress Test (Cont'd)




■ Three Individual Tests:

1. Daily – simulate daily production operations.
 - Instruments: FM1, FM2, FM3, FM4, and FM5
 - PGEs: 1.0P2, 1.1P8, 1.2P1, 1.4P1, 2.1P1, 2.2P1
2. Monthly – simulate single month reprocessing job.
 - Instruments: PFM, FM1, FM2, FM3, FM4 (not enough sample data for FM5, so use PFM as a proxy for FM5 load)
 - PGEs: 1.1P8, 1.2P1, 1.4P1, 1.4P2, 2.1P1, 2.2P1, 2.3P1, 2.3P2, 3.1P1
3. Large – simulate large six month reprocessing job.
 - Instruments: (same as monthly)
 - PGEs: (same as monthly)




AMI-P Stress Test (Cont'd)




- Job scheduling coordinated with custom multithreaded Perl server application and slightly modified version of Perl_Lib
- PGEs submitted to SGE via AJSS
- Cluster load and blade statistics monitored through Ganglia:
 - CPU Load
 - Physical Memory Utilization
 - Swap Memory Utilization
- CPU cores available for processing:
 - IBM Power 6 - 112 cores (108 available)
 - Intel x86 – 64 cores
- Disk I/O statistics were not available due to lack of GPFS statistics collection software

AMI-P Stress Test Results

PGE Jobs per Architecture




CPU Time per Architecture



- Most of the computation time was spent on running CER1.1P8 on the Power 6s



AMI-P Stress Test Results (Cont'd)





- Analysis focusing on the 6 month test results
- Workload heavily favored Power 6 platform
- CER1.1P8 ran for the most amount of time per instance
- Caveats:
 - CER1.4P2, CER2.3P2, and CER3.1P1 encountered environment misconfiguration errors in the stress tester causing them to fail to launch. The expected run times for these PGEs is greatly overshadowed by CER1.1P8
 - FLASHFlux processing ran alongside CERES stress testing
 - Unable to record disk I/O statistics due to lack of necessary software

PGE	Number of Instances	Average Runtime (minutes)	Architecture
CER1.1P8	915	203.65	Power 6
CER1.2P1	915	0.35	x86
CER1.4P1	905	0.34	x86
CER1.4P2	-	-	x86
CER2.1P1	8	0.24	Power 6
CER2.2P1	793	2.99	Power 6
CER2.3P1	22	0.15	Power 6
CER2.3P2	-	-	Power 6
CER3.1P1	-	-	Power 6


Total Wall Time: 30 Hours and 35 Minutes
Total CPU Time: 3155 Hours and 42 Minutes




AMI-P Stress Test Results (Cont'd)




AMI-P Stress Test Results (Cont'd)



AMI-P Stress Test Results (Cont'd)





AMI-P Stress Test Results (Cont'd)



1.1P8 x86 Memory Leak

CER1.1P8 Memory Utilization (bg102)






Backup Charts



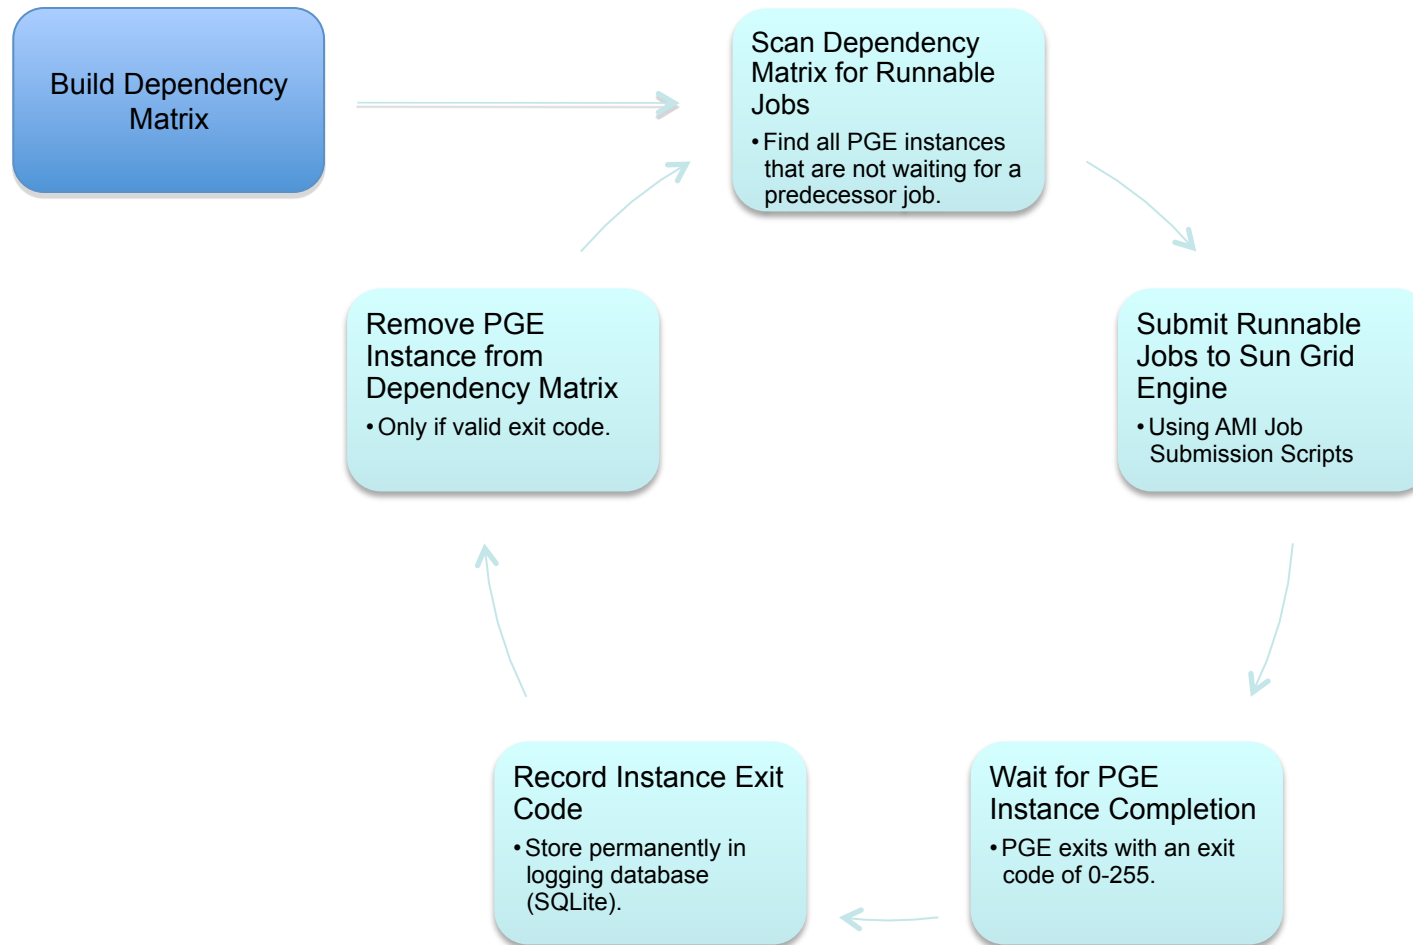
Notional Data Release Schedule



- Covers open + 7.5 months release Edition1-CV BDS, ES8, ES9 & ES4
 - 6 months observations
 - 1 month prepare gains & spectral response
 - 2 weeks produce and validate data


- Covers open + 13.5 months release Edition2 BDS, ES8, ES9 & ES4
 - 12 months observations
 - 1 month prepare gains & spectral response
 - 2 weeks produce and validate data

AMI-P Stress Test Control Flow



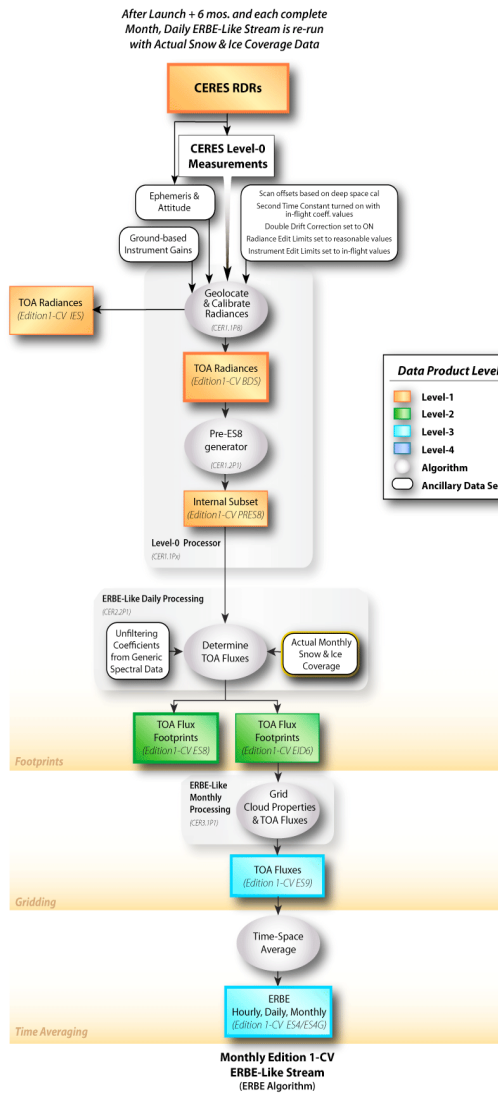


Data Sources




Type of Data	Source	Frequency / # of Files	Volume	Comments
NPP RDR	Land PEATE	131 – 134/day	186MB/day	In case of NPP, RDRs also contains spacecraft diary
VIIRS Sub-sampled Radiance and Geolocation	Land PEATE	288/day	60GB/day	CERES provided code to sub-sample at Land PEATE
VIIRS VAOT	Land PEATE	288/day	840MB/day	Produced by Land PEATE
SURFMAP (Snow/Ice)	NCAR	1/day	66 MB/mo	Used for Clouds and ERBElke Processing
SMOBA Ozone	NCEP	1/day	71MB/mo	
SURFMAP (Snow/Ice)	NOAA/NESDIS	1/day	630 MB/mo	Used for Clouds
MCIDAS Geostationary Data	University of Wisconsin	24/day/satellite	71GB/mo	5 geostationary satellites used per month
MATCH Aerosols	NCAR	2/day	164MB/mo	
Meteorological and Ozone data	GMAO	174 GB/mo	9/Day (CERESG5.2) 8/Day (CERESG5.4)	Currently receiving 2 streams G5.2 and G5.4

Data Processing Flow (Cont'd)





AMI Job Submission Scripts (AJSS)



- Leverages library routines from Perl_Lib.
- One Customized AJSS package per PGE.
- Verifies the following preconditions are true prior to executing a PGE instance:
 1. Environment variables are properly set.
 2. Required inputs exist.
 3. Expected outputs do not exist.
- Forwards the environment settings to Sun Grid Engine for remote job execution.